

First Hit   Fwd Refs

Generate Collection

Print

L3: Entry 16 of 29

File: USPT

Aug 15, 2000

DOCUMENT-IDENTIFIER: US 6105080 A

TITLE: Host adapter DMA controller with automated host reply capability

Brief Summary Text (9):

It is also generally known in the art to use scatter/gather DMA techniques to permit flexibility in the distribution of data to be exchanged between a source and destination in a DMA transfer. Such scatter/gather techniques enable a DMA to retrieve data for the transfer from non-contiguous memory locations in the source memory (gather) and/or to store the retrieved data in non-contiguous locations of the destination memory (scatter). As presently known in the art, scatter/gather DMA devices use at least one scatter/gather list that specifies to the DMA transfer circuits (DMA engine) a series of source blocks of data to be retrieved and a corresponding series of destination blocks into which the retrieved data is to be stored.

Brief Summary Text (11):

In processing of a host I/O request, the host adapter generally constructs a scatter/gather DMA block list to define the entire transfer required between the adapter's local memory and the host system's memory. At completion of processing of the specified scatter/gather list(s) the host adapter's processor is interrupted to complete the I2O specific transaction. For example, the host system may require a response from the I2O compatible host adapter to indicate that the requested transfer is complete and the status of the completion.

Brief Summary Text (16):

The present invention solves the above and other problems, thereby advancing the state of the useful arts, by providing a DMA design which integrates DMA transfers, as controlled by a scatter/gather list, with response transmission. The DMA controller of the present invention defines special markers (also referred to as flags or flag bits) associated with a DMA scatter/gather list which define a response to be transmitted to the requesting host system. As the DMA controller processes source and destination memory blocks defined by one or more scatter/gather lists, it monitors for the presence of such a marker flag signifying the need to transmit a response to the attached host system. When such a flag is sensed, the DMA next fetches information associated with the scatter/gather list required to send the desired response. The desired response is fetched and/or constructed and transmitted to the attached host system. The DMA controller then continues with any remaining portions of the DMA transfer defined by the scatter/gather list.

Brief Summary Text (20):

Preferably, a separate "end of list" flag is provided to indicate the end of the corresponding list and hence the end of the particular DMA transfer. The reply marker flag is therefore independent of the end of transfer flag/indication permitting a reply to be sent within a DMA transfer sequence as well as at the end of a transfer sequence. In particular, multiple such replies may be generated as required within a single DMA transfer sequence.

Detailed Description Text (6):

As is known in the art of DMA controllers, scatter/gather lists are useful to

describe a desired DMA transaction. Such a scatter/gather list defines the source and destination for a list of blocks which, together, comprise DMA transaction. DMA 100 in host adapter 106 is a DMA controller in accordance with the present invention which is adapted to perform the exchange of information between local memory 130 and host memory 122. DMA 100 uses at least one scatter/gather list to direct the processing of the desired DMA transaction. DMA 100 accesses local memory 130 via internal bus 152 and accesses host memory 122 indirectly via PCI bus 102.

Detailed Description Text (7):

For example, in processing a write I/O request, DMA 100 retrieves information from host memory 122 and moves it to local memory 130 for further processing within host adapter 106. More particularly, local memory 130 may be a cache memory in a host adapter 106 used for storage subsystem management (e.g., RAID storage subsystem management). The data to be written may be stored in host memory 122 in non-contiguous blocks. These locations are made known to host adapter 106 via a scatter/gather list provided by host system 104. In like manner, the locations in local memory 130 to which the blocks are to be written may be non-contiguous (scattered throughout the memory). Read operations operate in a similar manner but in reverse. Specifically, a host system 104 read I/O request will request retrieval of particular blocks or types of data. In the case of a storage subsystem host adapter 106, the requested data, if not already present, is retrieved from the storage devices (not shown) and temporarily saved in local memory 130 (i.e., cache memory). The retrieved data may be scattered about within local memory 130. The locations in host memory 122 to which the retrieved data is to be transferred may likewise be non-contiguous. These locations are made known to host adapter 106 via a scatter/gather list provided by host system 104.

Detailed Description Text (8):

The scattered locations within local memory 130 are known to DMA 100 via a scatter/gather list constructed within host adapter 106 by CPU 128.

Detailed Description Text (9):

DMA 100 is operable to exchange the requested data between local memory 130 of host adapter 106 and host memory 122. As is previously known in the art, CPU 128 may construct a composite scatter/gather list which joins the two supplied scatter/gather lists. However, so joining the scatter/gather lists requires processing overhead within host adapter 106. In particular, as is taught in the art, such a composite list is built such that each entry includes a source location, a corresponding destination location, and a block size corresponding to both. Since the scattered block in host memory 122 and local memory 130 may not be of equal size, CPU 128 must construct the composite list so as to define portions of source and destination blocks having equal sizes.

Detailed Description Text (13):

As noted further below, in the preferred embodiment, the reply message to be sent is pre-built by the processor of the host adapter which uses the DMA controller 100. The next entry in the S/G list following an entry having the reply bit set therefore includes a reference (e.g., pointer in memory or other identification means) to a pre-built reply message. This reference is preferably stored within the source address field of the next S/G list entry. Further, the destination address field within the next entry following the entry having a reply flag set defines the address to which the pre-built reply message is to be sent. Those skilled in the art will recognize that the reply message reference information, like the reply flag itself, could be stored in meta-structures associated with the scatter/gather list(s) rather than physically within the scatter/gather list(s). Where the DMA of the present invention is utilized with an I2O compatible host, it may be preferred that the reply flag and reference information be stored other than physically in the scatter/gather list. The I2O standards prescribe certain format standards for scatter/gather list. In storing the reply information elsewhere, the DMA of the present invention may be capable of directly utilizing the scatter/gather list